

SPECIFICATION

Electronic Version 1.2.8

Stylesheet Version 1.0

[METHOD AND APPARATUS FOR ETHERNET PRIORITIZED DEVICE CLOCK SYNCHRONIZATION]

Cross Reference to Related Applications

This patent application is being filed concurrently with commonly assigned U.S. Patent Application entitled, "A Method For Managing Bandwidth On An Ethernet Network", Serial No. ##/###,###, filed April 1, 2002 (Attorney Docket No. SAA-79-1 (401 P 282)); the content of which is expressly incorporated herein by reference. This patent application is related to U.S. Patent 6,223,626 entitled "SYSTEM FOR A MODULAR TERMINAL INPUT/OUTPUT INTERFACE FOR COOMUNICATING MESSAGE APPLICATION LAYER OVER ETHERNET TO TRANSPORT LAYER;" the content of which is expressly incorporated herein by reference. This patent application is related to and claims priority to U.S. Patent Application entitled "COMMUNICATION SYSTEM FOR A CONTROL SYSTEM OVER ETHERNET AND IP NETWORKS," Serial No. 09/623,869, filed September 6, 2000 (Attorney Docket No SAA-9); the content of which is expressly incorporated herein by reference.

Background of Invention

[0001] Technical Field. The present invention relates generally to synchronization of control system device clocks through a communication network. More specifically, the present invention relates to accurate Ethernet synchronization of device clocks using IEEE 802.1p protocol.

[0002] Background of Invention. The Institute for Electrical and Electronic Engineers (IEEE) defines the Ethernet standard as IEEE Standard 802. This standard establishes an

Ethernet network configuration guideline and specifies how elements of an Ethernet network interact. Network equipment and protocols can communicate efficiently when adhering to the IEEE standard.

[0003] Network protocols are standards that allow devices to communicate. A protocol defines how network devices identify one another. Another protocol defines the format of the transmitted data and how the data gets processed once it reaches its destination. Protocols such as transmission control protocol/internet protocol (TCP/IP) (for UNIX, Windows NT, Windows 95 and other platforms) also define procedures for handling lost or damaged transmissions or "packets".

[0004] Quality of service in the TCP/IP is made up of five layers, i.e. application, transport, network, data, and physical layers. The application layer is for application and user access, authorization, and encryption. The transport layer has TCP rate control and port access control. The network layer balances loads, reserves resources, and contains type of service bits, and path controls. The data link layer is where IEEE802.1p/Q frame prioritization takes place as well as logical port access control. Finally, the physical layer has bit error correction, physical security and port access. Ethernet is a shared media with rules for sending packets of data. These rules protect data integrity and avoid conflicts. Nodes determine when a network allows packets to be sent. In industrial control applications, it is possible for two nodes at different locations to send data concurrently. When both devices transfer a packet to the network concurrently, a collision will result.

[0005] Minimizing these collisions in factory automation applications is a critical portion of the design and operation of their networks. An increase of collisions in industrial control environments is frequently caused by increases of control-system devices on the network. This creates contention for network bandwidth and slows network performance.

[0006] Network quality of service is important for proper and predictable industrial control system performance. There are a number of factors that may diminish network performance. The first is delay, which is the time a packet takes to go from the sender to the receiver device via the network. Long delays put greater stress on the transport protocol to operate efficiently, particularly motion control, drives and robots

applications. Long delays imply large amounts of network data held in transit. Delays affect counters and timers associated with the protocol. In the TCP protocol, the sender's transmission speed is modified to mirror the flow of signal traffic returning from the receiver, via the reply acknowledgments (ACK's) that verify a proper reception. Large delays from senders and receivers make the feedback loop insensitive. Delays result in the protocol becoming insensitive to dramatic short-term differences in industrial control system network load. Delays affecting interactive voice and video applications cause systems to appear unresponsive.

[0007] Jitter is another network transit delay. Large amounts of jitter cause the TCP protocol to conservatively estimate the round trip message time. This creates inefficient factory automation protocol operation by requiring timeouts to reestablish the flow of data. A large quantity of jitter in user datagram protocol (UDP) based real time applications such as an audio or video signal is intolerable. Jitter creates distortion in the signal, which then must be cured by enlarging the receiver's reassembly playback queue. The longer queue delays the signal, making interactive communication difficult to maintain, detrimentally for factory automation.

[0008] A third network issue is its bandwidth, the maximal industrial control data transfer rate. Bandwidth may be restricted by other traffic that shares common elements of the route as well as the physical infrastructure limitations of the traffic path within the factory automation transit network.

[0009] Reliability is yet another industrial control system concern. Network reliability may be measured by the average error rate of the medium. Improperly configured and poorly performing switching systems may change the order of packets transmitted. Packets may be delivered to the receiver in a different order than what was sent. Packets could possibly be dropped while in transient routing loops. Industrial control network transit paths that are not reliable or that are prone to error may also require the lost packets to be sent another time. TCP does not differentiate between messages lost because of packet corruption and messages lost because of network congestion. The sender is subject to the same congestion avoidance behavior responses when packets are lost. The industrial control network reacts by reducing message rates even more and forcing congestion avoidance algorithms despite the

fact that no actual congestion was experienced on the network. Poor reliability in voice and video UDP applications causes distortion in the original analog signal at the receiver's end.

[0010] Ethernet's general lack of message prioritization and the openness of the TCP/IP protocol may introduce latent performance flaws relating to network traffic. Industrial control network traffic bursts can result in message losses and slow responses caused by non-critical network traffic. Categorizing traffic may be implemented to ensure that critical factory automation traffic will always flow despite the demands of less important applications. The prioritization of industrial control network traffic enables predictable performance for the most critical application traffic.

[0011] IEEE802.1p standard defines how network frames are tagged with user priority levels ranging from 7 highest to 0 lowest priority. IEEE802.1p compliant network infrastructure devices, such as switches and routers, prioritize network traffic delivery according to the user priority tag. Higher priority tagged frames are given precedence over lower priority or non-tagged frames. This means that time critical data receives preferential treatment over data that is not considered time critical.

[0012] Quality of service mechanisms can be incorporated at any or all of the five layers of the TCP/IP stack and the positioning of the key quality of service mechanisms. Some of these mechanisms are inherent in the protocols rather than being explicitly added on for quality of service control.

[0013] Ideally, networks provide services to an unlimited number of users, at any number of locations, without any limitations on what can be transported. Industrial control systems should have well-specified delays in service access or usage, with contained packet loss rates, and without any congestion. The grade of service for a specific traffic flow is decided by the user/system that is generating the flow. Service quality is also influenced by enterprise administrators who control the budget. Network managers in charge of one of the network components are another factor. Even a service provider that supplies the WAN facilities influences service. By defining a small number of classes of service, the network allows for packet differentiation, maintains an acceptable quality of service, and yet still provides the services at reasonable cost and complexity.

[0014] Accurate clock synchronization is particularly critical in industrial control applications where multiple network components request proper timing. Fast synchronization of all device clocks is required through networks, with preferably, a deviation of less than a few hundred microseconds. At initialization, devices start their clocks when prompted and therefore their times differ from one another. The desire is for all initializations to occur as close to simultaneous as possible. When timing is not concurrent between multiple industrial control devices on the same network, real time data exchange quality of service suffers.

[0015] Motor control, drives, robots, factory automation, and electrical distribution are just a few of the potential applications for industrial controls linked with Ethernet. Minimizing the timing differences between industrial control device clocks is much more important than in most Ethernet applications and therefore is a desired goal.

Summary of Invention

[0016] It is an object of the invention to provide an Ethernet industrial control system for transferring messages, wherein the messages are tagged with identifiers of varying levels of priority, ranging from a highest priority to a lowest priority.

[0017] In accordance with the invention, the system comprises a serial network bus and a means for placing a message having a higher priority identifier onto the bus before placing a message with a lesser priority identifier onto the serial network bus.

[0018] Other features and advantages of the present invention will be apparent from the following specification taken in conjunction with the following drawings.

Brief Description of Drawings

[0019] FIGURE 1 is a simplified block diagram of a network with n devices;

[0020] FIGURE 2 is a simplified block diagram of network exchanges and code execution;

[0021] FIGURE 3 is a simplified block diagram of network device event traffic;

[0022] FIGURE 4 is a block diagram of asynchronous transfer mode network time stamping;

- [0023] FIGURE 5 is a simplified block diagram of prioritized network traffic classes using dedicated parallel communication stacks;
- [0024] FIGURE 6 is a simplified block diagram of Ethernet reduced stack clock synchronization;
- [0025] FIGURE 7 is a simplified block diagram of classic Ethernet traffic throughput;
- [0026] FIGURE 8 is a simplified block diagram of Ethernet architecture with multiple stacks and prioritized frames;
- [0027] FIGURE 9 is a diagram of the Ethernet header fields;
- [0028] FIGURE 10 is a simplified diagram of the interrelationship of tagged and untagged device traffic;
- [0029] FIGURE 11 is a simplified block diagram of tagged and untagged device traffic;
- [0030] FIGURE 12 is a simplified block diagram of tagged device to tagged device traffic;
- [0031] FIGURE 13 is a simplified block diagram of Ethernet high priority, reduced stack, clock synchronization; and,
- [0032] FIGURE 14 is a simplified block diagram of a sample network.

Detailed Description

[0033] While this invention is susceptible of embodiments in many different forms, there is shown in the drawings and will herein be described in detail a preferred embodiment of the present invention with the understanding that the present disclosure is to be considered as an exemplification of the principles of the invention and is not intended to limit the broad aspect of the present invention to the embodiment illustrated.

[0034] The quality of service characteristics of an industrial control network can be managed using mechanisms operating at the edge of the network or within its core. Quality of service may be controlled by reserving a fixed amount of bandwidth for mission-critical applications or preventing specific users from accessing restricted data like WWW destinations.

- [0035] Assigning a higher priority to traffic to and from specific customers, limiting the bandwidth that can be consumed by voice over IP traffic, or designating which types of traffic may be dropped when congestion occurs are other quality of service controls. End-to-end solutions include regulating individual or traffic flows, processing quality of service information within the network, and monitoring the configuration of the network.
- [0036] At initialization, devices start their clock and so, they differ from each other. Real time data exchanges like clock synchronization, global data or IO scanning should be deterministic and should have guaranteed time delivery, whatever other traffic is on the network.
- [0037] Class based queuing allows traffic to be divided into a hierarchy of classes that are determined by quality of service attributes. IEEE802 provides for 8 levels of message priority. Class based queuing also allows borrowing of bandwidth among the classes. Queuing mechanisms can be pre-configured. The traffic flow through an Ethernet stack architecture can be implemented through a policy-based management system, or can be controlled via user-driven signaling protocols. Ethernet strikes a balance between speed, cost and ease of installation. Preferably, clock synchronization frames are tagged with the highest network priority level.
- [0038] Prioritization of traffic on an industrial control local area network (LAN) allows synchronization requirements to be signaled to LAN switches and routers. The IEEE802.1p specification defines three bits within the IEEE802.1Q field in the MAC header, a part of the open systems interconnection (OSI) Layer 2. The IEEE802.1Q field is designed to support virtual-LAN (VLAN) interoperability and to support traffic priorities. VLAN allows the use of multiple, logical, virtual Ethernets over a single physical Ethernet. Three bits support up to 8 settings, used for classes of traffic and priorities.
- [0039] A reduced UDP/IP communication stack enables messages to be put on a serial network bus faster than ordinary communication stacks. Device clock synchronization messages preferably receive the highest priority level tags and additionally are transmitted through a reduced UDP/IP communication stack, optimizing clock synchronization. Implementation of message priority tagging and the use of a reduced

UDP/IP communication stack for clock synchronization creates improved accuracy of all network device clocks, preferably with an accuracy of less than a few hundred microseconds.

[0040] Referring to FIGURE 1, an Ethernet based industrial control system 10 for properly synchronizing local device clocks 20, 22 to the clock manager 18 is illustrated. FIGURE 1 is a block diagram of one embodiment of the present invention. There is fast synchronization of all device clocks 20, 22, through the network 24, with accuracy less than a few hundred microseconds.

[0041] Referring to FIGURE 2, network exchanges 32, 34 and 36 and user code executions 38, 40, and 42 are synchronously scheduled. This diagrams the synchronous scheduling of network exchanges 32, 34 and 36 and of code execution 38, 40, and 42 where all devices have consistent data when starting code execution and distributed control capacity 44 is provided.

[0042] Referring to FIGURE 3, a diagram of accurate clock synchronization also providing capability for fast periodic tasks. Periods T as low as 2 ms are possible for all n devices on the network. There is event 56 time stamping capability in each device 52, 54, with discrimination of events distributed in the overall system and separated by only 1 ms, shown as $\text{time} + 1 \text{ ms}$ at device number n .

[0043] Referring to FIGURE 4, a diagram of clock synchronization and quality of service with an asynchronous transfer mode (ATM) network environment. The hardware components 62 provide precise time stamping of sent and received frames. With the network of FIGURE 4, there is preferably an 80 microsecond accuracy achieved with the quality of service native to the ATM network 68.

[0044] FIGURES 5 and 6 refer to the present invention, using Ethernet with IEEE802.1p prioritization. Clock synchronization frames are preferably tagged with the highest priority and are put on the network through a reduced stack UDP/IP. The use of a reduced stack and frames tagged with priority aids efficient device clock synchronization. Lower priority traffic must wait for higher priority frames to be transmitted. FIGURE 5 is a diagram an embodiment of synchronization and quality of service with Ethernet. The Ethernet standard 802.1p is used to prioritize frames and

traffic is sent through the 802.1p switch 72. The priorities are managed from application to application inside devices. In this embodiment, the highest priority level is given to clock synchronization frames 76.

[0045] FIGURE 6 refers a diagram of an embodiment of clock synchronization with Ethernet 82. Device clock 88 synchronization frames get prioritized and pass through a reduced stack. Other Ethernet traffic passes through standard protocol channels, onto the network. The embodiment is developed with a standard hardware device with no additional hardware component. Software modifications are made to manage clock synchronization by implementation of IEEE802.1p priorities on the network and in the device.

[0046] FIGURE 7 is a diagram of the classic embodiment of clock synchronization with Ethernet 92. This embodiment does not employ IEEE802.1p nor use a reduced stack. Frames are sent in random order with synchronization frames possibly waiting T times the number of frames 94. The time 94 needed for traffic to pass through the TCP/IP 96 can be contrasted with the Ethernet architecture represented in FIGURE 8.

[0047] Referring to FIGURE 8, in an embodiment of the present invention wherein a reduced stack (UDP/IP) 102 is allocated to the highest priority frames, clock synchronization 104. A second reduced stack (TCP/IP) 106 addresses slightly lower prioritized frames, such as I/O scan frames 108. The last stack (TCP/UDP/IP) 110 handles the lowest priority traffic, such as FTP 112 and HTTP 114 frames. This embodiment does employ IEEE802.1p and reduced stack Ethernet 116 architecture. Frames are sent in the priority order and clock synchronization frames are only delayed by t (time passing through the reduced UDP-IP stack 102).

[0048] Referring to FIGURE 9, Ethernet message headers are shown to represent an untagged header 122 without IEEE802.1p and a tagged header 124 with IEEE802.1p tag control information including the 3-bit priority field 126. FIGURE 9 diagrams the IEEE802.1p protocol. Four additional bytes are inserted in the Ethernet layer 2 MAC frames as tag control information with a 3-bit field for priority.

[0049] FIGURES 10, 11, and 12 display the interaction compatibility of untagged and tagged devices and frames. FIGURE 10 is a diagram depicting the compatibility

between standard Ethernet 136 and IEEE802.1p 132 frames. The switch 134 ensures the compatibility between a tagged and an untagged device. The switch 134 adds the VLAN information to untagged frames by assigning them a priority of zero, the lowest priority. The switch 134 also removes the VLAN information from frames addressed to untagged devices.

[0050] FIGURE 11 is a diagram depicting the interrelationship between a tagged 142 and an untagged 144 device. In this embodiment, frames tagged by a programmable logic controller 142 arrive on port 3 and are automatically untagged by the switch 146 and sent to port 10. On the return path, frames arrive at port 10 and are tagged with a VLAN of one and a priority of zero. Frames are then switched to port 3 where they are transmitted out of the switch tagged.

[0051] FIGURE 12 is a diagram depicting an interrelationship between two tagged devices 152 and 154. In this embodiment, frames tagged by programmable logic controllers 152 arrive on port 3 and 10, the switch 156 detects them as tagged, and forwards them as configured. Frames are tagged with a priority level. Frames are switched corresponding to this priority.

[0052] Referring to FIGURE 13, a diagram depicting the results of a complete transaction response time from application to application, request and response. In classical stack architecture, there is no guaranteed time synchronization. A classical Ethernet stack can be congested and deterministic exchanges become impossible. With a reduced stack and IEEE802.1p as shown, time synchronization is guaranteed. Fast synchronization of all device clocks is achieved with accuracy 162 less than 500 μ s, in the worst case.

[0053] Turning to FIGURE 14, shows a diagram depicting an embodiment of IEEE802.1p implementation over a complete architecture. Network devices can include, but are not limited to, PLCs 178, 180, and 182, an IEEE802.1p switch 176, some type of human/machine interface 184, and audio and video 172 and 174 devices.

[0054] Potential applications using the preferred embodiment of the present invention include motion control, drives and robots application requiring fast synchronization, electrical distribution applications requiring discrimination of events, automation

